

**IN THE UNITED STATES DISTRICT COURT  
FOR THE EASTERN DISTRICT OF TEXAS, MARSHALL DIVISION**

FUNIMATION ENTERTAINMENT  
A TEXAS LIMITED PARTNERSHIP,  
Plaintiff,

V.

DOES 1 – 1,427,  
Defendants.

Case No. 2:11-cv-00269

Judge: Hon. David Folsom

## DECLARATION OF SETH SCHOEN

I, Seth Schoen, declare as follows:

1. I am a Senior Staff Technologist with the Electronic Frontier Foundation (EFF), and I make this declaration on my own personal knowledge. I have worked with computers and computer networks for over a decade, have testified about electronic communications systems in two courts and before the United States Sentencing Commission, and have submitted declarations similar to my present declaration to the Federal courts in at least eight other matters.

2. The purpose of this declaration is twofold. The first is to set forth facts, which were readily available to Plaintiff from free, public Internet sources at and before the time it filed suit, that establish that many of the unnamed Defendants in the above-referenced case (hereinafter “Does” or “Doe Defendants”) use Internet connections almost certainly located physically outside of the State of Texas. The second purpose of this declaration is to respond to assertions made by Plaintiff that might give a misleading impression of how unique BitTorrent is or how likely it is that various Defendants interacted with each other or were aware of each other in the course of uploading or downloading the motion picture whose copyright Plaintiff accuses them of infringing.

**STATEMENTS RELATING TO PERSONAL JURISDICTION**

3. By reviewing Exhibit A to the Complaint, I compiled a list of the Internet Protocol (IP) addresses that Plaintiff attributes to each of the Doe Defendants.

4. There are many tools freely available to the public that help reveal where a person using a particular IP address is likely to be physically located. This process is often referred to as “geolocation.” This information is commonly used for many purposes, such as customizing the language or content of web sites based on inferences about where visitors are accessing the site from. For example, Google, Inc., uses geolocation to choose to display its web site in German to people coming from Germany, in French to people coming from France, and so on. It also uses geolocation to display ads and results related to particular cities or regions to people accessing its site from those cities or regions.

5. One means of learning about where an IP address is physically located is known as “reverse domain name service lookup” or “reverse DNS.” When an Internet service provider (“ISP”) allocates or prepares to allocate IP addresses to customers, it typically creates and publishes database records assigning a human-readable “domain name” to each numerical IP address. The reverse lookup information can be obtained by anyone using a program such as “host,” which is a standard program included with many computer operating systems, or with any of several web-based tools such as the DNS lookup service at <http://lookupserver.com/>.

6. One of the purposes of reverse DNS is to help interested parties learn more about what a computer is used for, what organization’s network it is connected to, and, in many cases, where the computer is physically located. Typically, for home users of dial-up or broadband connections, such as DSL or cable-modem services, a domain name obtained from reverse DNS will identify which ISP assigned the IP address.

7. In addition, such a domain name will frequently incorporate an approximate physical location, such as the name of a municipal area, state, or region. For example, Doe #17 is identified by the IP address 69.136.223.171 and described by Plaintiff as a subscriber of Comcast Cable. The reverse DNS database identifies this computer as c-69-136-223-

171.hsd1.pa.comcast.net, confirming Plaintiff's suggestion that Doe #17 is a Comcast ("comcast.net") customer, likely using Comcast's cable Internet service, but adding the additional detail that the likely physical location of the computer is in or near Pennsylvania ("pa"). This means that in all likelihood, the individual who used this IP address is located in the State of Pennsylvania.

8. Although Internet service providers are not required to publish this information, and although it is sometimes only given to state-level precision, it can, when available, be a useful source of data about where an individual Internet connection is most likely located.

9. For each of the 1,427 IP address that were references in this suit, I used the "host" program to perform a reverse lookup against the publicly-accessible reverse DNS service.

10. The results of this process generally confirmed Plaintiff's association of particular IP addresses with particular ISPs. Additionally, the results of this process generally strongly suggested a geographic location for most individual defendants. In other words, most of the Does listed in this lawsuit can be associated by the host reverse DNS look-up with both an Internet service provider and a geographic location.

11. Reverse DNS records indicate that Does in this lawsuit include customers with Internet connections located in virtually all areas of the United States, including some in or near Hawaii, Washington, D.C., Pennsylvania, Georgia, Kansas, Florida, and other states and regions throughout the United States.

12. In addition to reverse DNS information, another means of learning where an IP address is located is to use a public database operated by the American Registry for Internet Numbers ("ARIN"). ARIN is the authority responsible for the initial allocation of IP addresses to ISPs located in the United States. ARIN maintains public records indicating to whom a given IP address has been allocated. Large ISPs may apply to ARIN multiple times to receive multiple "blocks" or ranges of IP addresses. Each such block may be dedicated to a particular purpose or geographic area.

13. The ARIN database can be searched using a public web site provided by ARIN at

<https://www.arin.net>, or by using a program called “whois,” which is a standard part of some operating systems and performs the same database-searching function. There is no charge for searching the ARIN database.

14. For example, Doe Defendant #202 is identified by the IP address 71.239.193.77. I searched the ARIN whois database for this address and learned that this address is part of a network assigned to “Comcast Cable Communications.” The whois record also a field called “Name” which has the value “ILLINOIS-29.” This information is readily available at <http://whois.arin.net/rest/net/NET-71-239-0-0-1/pft>. Combined with the reverse DNS record for this IP address, which is c-71-239-193-77.hsd1.il.comcast.net, there is a strong inference that the user of this IP address resides in or around Illinois, not Texas.

15. In addition, several companies collect and continually update geographic information about IP address locations from a variety of data sources, and collect this information in databases called “geolocation databases.” Geolocation databases are commonly used by web site operators who are interested in finding out the approximate physical location of their web visitors. Since web site operators are often very interested in such information, there is considerable demand for geolocation databases.

16. Geolocation databases may be sold or given away for free. One very popular geolocation database is the “GeoIP” database maintained by MaxMind, Inc., a Boston company that specializes in geolocation technology. In addition to other sources of information, MaxMind explains that it “employ[s] user-entered location data from sites that ask web visitors to provide their geographic location” in order to learn which IP address ranges correspond to which cities and states. MaxMind, <http://www.maxmind.com/app/ip-locate> (last visited July 7, 2011).

17. A version of the MaxMind GeoIP geolocation database is freely available for anyone to download from MaxMind. The company claims that this free version can determine the location of “79% [of U.S. IP addresses] within a 25 mile radius.” MaxMind, <http://www.maxmind.com/app/geolitecity> (last visited July 7, 2011).

18. I downloaded this freely available database and looked up each mentioned IP address

in it, obtaining an estimated city and state location for each such address.

19. Because DSL and cable modem connections are provided from local hubs to users in a particular geographic region, there is good reason to believe that the geographic location data obtained by these methods actually reflects the physical location of the Internet connection, at least in general terms. In other words, although geolocation data is not perfectly accurate, the geographic designations obtained by these methods likely indicate the approximate locations of the residences or other venues where the Does use their Internet-connected computers.

20. I have attached hereto as Exhibit A to this Declaration a list of the reverse DNS names of the Doe Defendants' distinct IP addresses, as well as the estimated physical location of each such IP address according to the freely available version of the MaxMind GeoLite City database.

21. In my experience, computer professionals are generally aware of the existence and function of the reverse DNS and whois services, as well as geolocation databases such as the GeoIP database, and would use any or all of these sources of information when they needed to learn where a given IP address was physically located. These techniques are readily and easily available to Plaintiffs, their attorney, and to the computer professionals they have employed to perform the investigations leading to this lawsuit.

22. Though the MaxMind GeoLite City database and reverse DNS records are not perfectly accurate, I know of no reason to think that either source of information has a bias that makes it more or less likely that an individual IP address will appear to be located in Texas.

23. From the information available from the MaxMind geolocation database, 114 (one hundred fourteen) of the IP addresses appear to be located in the State of Texas, 1248 outside of Texas, and 65 are not assigned to any location by the database. This puts around 8% of the IP addresses in the State of Texas, comparable to the 8.1% of the population of the United States as a whole that resides in Texas according to the 2010 Census.

24. Separately from the question of where Does reside, Plaintiff did not submit all the details of the investigations that led it to accuse these Does of copyright infringement. These

details could be important because simple methods of attempting to locate copyright infringers can easily go awry. For example, in 2008 researchers from the University of Washington found that, given then-prevalent methods for investigating BitTorrent transfers, it was straightforward to frame particular IP addresses for downloading files that they had not, in fact, ever attempted to download. The researchers experimentally framed their own laser printer and succeeded in eliciting false allegations of copyright infringement against it. See Michael Piatek, Tadayoshi Kohno, and Arvind Krishnamurthy, “Challenges and Directions for Monitoring P2P File Sharing Networks, or, Why My Printer Received a DMCA Takedown Notice,” in *Proceedings of the 3<sup>rd</sup> USENIX Workshop on Hot Topics in Security*, July 29, 2008, available at [http://www.usenix.org/event/hotsec08/tech/full\\_papers/piatek/piatek.pdf](http://www.usenix.org/event/hotsec08/tech/full_papers/piatek/piatek.pdf).

#### **STATEMENTS RELATING TO MASS JOINDER**

25. I reviewed the Complaint as well as some of the academic research on BitTorrent, as indicated below.

26. This Declaration responds to assertions made by the Plaintiff that might give a misleading impression of how unique BitTorrent is or how likely it is that various Defendants interacted with each other or were aware of each other in the course of uploading or downloading the motion picture whose copyright Plaintiff accuses them of infringing.

27. Plaintiff claims that BitTorrent is “unique in that it is designed to be used by a *group* of people working together.” Compl. ¶ 4.

28. Ultimately, however, BitTorrent is actually strikingly similar to file sharing systems that were at issue in previous litigation about peer-to-peer file sharing, and to the extent it is different, the differences result in less direct communication among users of the technology, not more.

29. First, BitTorrent is not the only system that has a swarming or multi-source download feature in which users can download simultaneously from several other users. Although this design was not a part of the earliest popular peer to peer systems such as Napster, it subsequently became quite widespread. For instance, the Kazaa and Gnutella software that was at issue in

several copyright infringement actions have a swarming download feature that works similarly to BitTorrent's. *See, e.g.*, L. Jean Camp, "Peer to Peer Systems", in Hossein Bidgoli (ed.), *The Internet Encyclopedia* (Wiley, 2004), vol. 3, at 30. ("In order to increase the speed of downloads and distribute the load on peer-provid[ed] files Limewire uses swarming transfers. Swarm downloading entails downloading different elements of files available on multiple low-bandwidth connections to obtain the equivalent service of a single broadband connection."); *see also* Alex Jantunen *et al.*, "Peer to Peer Analysis: State of the Art" (Tampere University of Technology, 2006) (noting that swarming supporting protocols include at least FastTrack, Gnutella, ED2K/Overnet and BitTorrent).

30. Second, BitTorrent's file-focused distribution provides users with *less* ability to identify and communicate with the peers with whom they exchange files than other technologies do. For example, Napster and KaZaA, unlike BitTorrent, referred to each user by a human-intelligible and somewhat memorable screen name, instead of a number. Napster and KaZaA have also offered users the ability to chat with one another. BitTorrent does not offer these features. There is no easy way for the various BitTorrent users who have uploaded or downloaded parts of a file to recognize, name, or communicate with one another.

31. While BitTorrent client software, like other peer-to-peer file sharing software, may provide a way for a user to view the IP addresses of peers, users are not required to do so in order to use BitTorrent. They do not have to select peers' IP addresses, because the selection of peers is done automatically. Indeed, since BitTorrent automates so much of the download process, many users likely do not even know how BitTorrent works. Most BitTorrent users have no reason to know how many or which other peers they might have communicated with in the course of downloading a file, or which addresses transmitted which portions of the file.

32. For example, the main screen of the popular Azureus BitTorrent software shows only a progress bar for the download, indication the percentage of the download that is complete, without mentioning other any other peers or their Internet addresses. *See, e.g.*, <http://torrent-search.us/images/torrent-clients/azureus-screenshot.jpg> (screenshot of Azureus

software in the midst of a download). Although interested users can learn about the role of peers or view their IP addresses, they are not required to do this.

33. I do not believe Plaintiff could have obtained direct evidence that any particular defendant shared portions of the copyrighted work at issue here with any particular other defendant, since BitTorrent does not provide a means for third parties to learn directly who is downloading files from whom.

34. Moreover, the plausibility that a given user downloaded a part of a file from any other particular user rapidly evaporates as the number of users becomes larger or as the users use BitTorrent at widely separated times. Both are true in this case. The number of users sued together in this case is over a thousand and, according to the records submitted by Plaintiff, they allegedly used BitTorrent at different times over the course of two weeks.

35. Both of these facts — the number of individuals named together and the different times of their alleged use of BitTorrent — make it highly implausible that all of 1,427 individuals sued jointly here uploaded or downloaded a part of the file from each other.

36. As to the different times for download specifically, the various Defendants are alleged to have used BitTorrent to transfer the movie file at very different times over the course of two weeks, which makes it even less plausible that they all could have communicated with one another. Exhibit A to the Complaint shows allegations of infringement on dates ranging from May 13, 2011 through May 27, 2011. Consistent with academic research on file-sharing using BitTorrent described below, this shows another reason why many individual defendants would never have communicated with other defendants: although some BitTorrent users may continue to share a file for a period of time after their download has completed, most do not.

37. Empirical research shows that most BitTorrent users do not remain connected for very long after their downloads are complete. These statistics can be measured by means quite similar to the techniques employed by Plaintiff's experts here. One large study observed that only 3.1% of BitTorrent users stayed connected (to upload to others) more than ten hours after their downloads completed; only 0.34% stayed connected over 100 hours. J. A. Pouwelse, P.



Garbacki, D. H. J. Epema, and H. J. Sips, *The BitTorrent P2P File-Sharing System: Measurement and Analysis* at 4, in Proceedings of the 4<sup>th</sup> International Workshop on Peer-to-Peer Systems, *available at* <http://www.springerlink.com/content/1251rj12233u051>.

38. Another study found that over 90% of users who successfully downloaded a file remained connected for less than a single day, while many users who attempted to download the file gave up entirely and disconnected within the first few hours. M. Izal, G. Urvoy-Keller, E. W. Biersack, P. A. Felber, A. Al Hamra, and L. Garcés-Erice, *Dissecting BitTorrent: Five Months in a Torrent's Lifetime* at 7, in Proceedings of the 5th International Workshop on Passive and Active Network Management Proceedings of the 4th International Workshop on Peer-to-Peer Systems, *available at* <http://www.springerlink.com/content/fg8hqw4136t0vtx9>.

39. Thus, it is highly unlikely all or even a significant number of the defendants who downloaded the subject copyrighted work here stayed on the network and became a source for another later-connecting defendant to download from days or weeks later.

40. Plaintiff states that each defendant “download[ed] Plaintiff’s film, piece by piece, from the other members in the group simultaneously.” Compl. ¶ 9. This statement could create two misconceptions about how BitTorrent works. In fact, a downloader receives a given “piece” of the file from only one other user, not from all “users who have already downloaded that piece.” BitTorrent does not permit downloading a particular piece of a file from more than one user at a time, although different pieces of the file can be downloaded from different users. Also, a downloader only communicates with *some* of the users in a *limited*, gradually changing “peer set” of generally no more than 50 peers. While it is plausible that *some* Doe Defendants shared *some* pieces of the allegedly infringing file with *some* of the other Defendants, Plaintiff’s assertion that “Defendants worked in concert” with the others is unsupported by its factual allegations or the nature of the BitTorrent protocol. Compl. ¶ 4.

41. As I stated earlier, many other modern P2P systems do support swarming downloads akin to BitTorrent’s, so it is hard to be confident that an infringer basically copied a work from one other user in the incidents at issue in prior file sharing litigation.

42. In any case, in all peer-to-peer file sharing networks, particular files can become more widespread throughout the network over time as new users obtain them from earlier users. Indeed, researchers have been able to quantify and analyze the spread of particular files in particular networks over time. Regardless of whether particular acts of copying involve two users or a greater number of users, the availability of a file logically depends on the decision of its original distributor to make it available. So the “events involving” people who share a file in a file-sharing system are equally “logically related” (or unrelated) in this sense, regardless of what technology underlies the file-sharing system.

I declare under penalty of perjury under the laws of the State of California that the foregoing is true and correct and that this document was executed in San Francisco, California.

Dated: January 30, 2012

By: \_\_\_\_\_  
SETH SCHOEN